




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : G06F 12/00	A1	(11) International Publication Number: WO 95/24682 (43) International Publication Date: 14 September 1995 (14.09.95)
(21) International Application Number: PCT/US95/02853		(81) Designated States: JP, KR, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).
(22) International Filing Date: 7 March 1995 (07.03.95)		
(30) Priority Data: 08/206,959 7 March 1994 (07.03.94)	US	Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>
(71) Applicant: SILICON GRAPHICS, INC. [US/US]; 2011 N. Shoreline Boulevard, Mountain View, CA 94039 (US).		
(72) Inventors: HANNAH, Marc, R.; 231 Covington Road, Los Altos, CA 94024 (US). NAGY, Michael, B.; 7516 Corrinne Street, San Ramon, CA 94583 (US).		
(74) Agents: SCHELLER, James, C. et al.; Blakely, Sokoloff, Taylor & Zafman, 7th floor, 12400 Wilshire Boulevard, Los Angeles, CA 90025-1026 (US).		

(54) Title: INTEGRATING TEXTURE MEMORY AND INTERPOLATION LOGIC



Best Available Copy

(57) Abstract

In a computer graphics system, a semiconductor chip (200) used in performing texture mapping. Textures are input to the semiconductor chip. These textures are stored in a main memory (205). Cache memory (208) is used to accelerate the reading and writing of texels. A memory controller (203) controls the data transfers between the main memory (205) and the cache memory (208). Also included within the same semiconductor chip is an interpolator (209). The interpolator (209) produces an output texel by interpolating from textures stored in memory (205). The interpolated texel value is output by the semiconductor chip, thereby minimizing transmission bandwidth as well as redundant storage of texture maps in a multi-processor environment.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

Integrating Texture Memory and Interpolation Logic

FIELD OF THE INVENTION

The present invention pertains to the field of computer display systems. More particularly, the present invention relates to an integrated texture memory and interpolation logic device used in performing texture mapping in a computer system.

BACKGROUND OF THE INVENTION

Computer systems are commonly used for displaying graphical objects on a display screen. These graphical objects include points, lines, polygons, and three dimensional solid objects. By utilizing texture mapping techniques, color and other details can be applied to areas and surfaces of these objects. In texture mapping, a pattern image or texture map is combined with an area or surface of an object to produce a modified object with the added detail of the texture map. For example, given the outline of a featureless cube and a texture map defining a wood grain pattern, texture mapping techniques can be used to map the wood grain pattern onto the cube. The resulting display is that of a cube that appears to be made of wood. In another

example, vegetation and trees can be added by texture mapping to an otherwise barren terrain model. Labels can be applied on computer-modeled package designs for visually conveying the appearance of an actual product. Furthermore, textures mapped onto geometric surfaces provide additional motion and spatial cues that surface shading alone might not be capable of providing. For example, a sphere rotating about its center appears static until an irregular texture or pattern is affixed to its surface.

Texture mapping involves using a texture having a function defined in texture space. The texture is then warped or mapped into an object space. Typically, a two-dimensional texture or pattern image is mapped onto a three-dimensional surface. This can be accomplished by associating a two-dimensional index (S,T) into an array of texture pixels (texels) in reference to points on the three dimensional surface. As the surface is rendered, the S and T values are interpolated and used to look up a texture value for each rendered pixel. The interpolated S and T values are generally not integer values. Consequently, these values often fall between the texture samples represented by the array of texels.

Several options exist for selecting a texture value, given real S and T values. One of the simplest options is to round S and

To the nearest integers and then select the texel corresponding to the rounded integer values. A more accurate representation is produced by interpolating between the four nearest samples that surround the real (S,T) location. Sometimes, a bi-linear interpolation algorithm is used. In other instances, higher-order interpolation algorithms are used for better results.

One potential problem with these unsophisticated selection processes described above is that, if the surface being texture-mapped is far from the viewpoint, severe aliasing of the texture may occur. This undesirable aliasing occurs because the interpolated (S,T) values may skip over large areas of the texture. A prior art technique, referred to as MIP-mapping, address this problem by precomputing multiple, filtered copies of the texture at successively lower resolutions. For example, a 256x256 texel array would be filtered and resampled to obtain maps at 128x128, 64x64, 32x32, 16x16, 8x8, 4x4, and 2x2 resolutions. The particular size of the texel array is chosen based on a computer parameter known as a level of detail (LOD). The LOD represents the relative distances between the interpolated (S,T) values. Each map size represents an integer LOD, and the computed LOD values are real numbers. High quality texture mapping is obtained by performing bi-linear interpolation in the map representing the integer LOD immediately above and below

the computed LOD of each pixel. Next, a linear interpolation is performed between the integer LOD's to obtain the texture value at the non-integer LOD. This process is known as tri-linear MIP-mapping.

Another problem is that in high performance computer graphics systems, the rendering operations are typically distributed among a number of processors. If this parallel organization were used to perform texture mapping, each processor would likely contain a copy of the entire texture map in its associated memory. This multiple storage of the entire texture map is both redundant and inefficient.

Yet another problem is that the transfer of data between the various processors and the memory consumes a great deal of the limited bandwidth. Allocating valuable bandwidth for this function tends to slow down the exchange of other needed information. Hence, the overall effect is that this imposes a heavy burden on the texture mapping process.

Thus, there is a need in prior art computer graphics systems for a fast, accurate, and efficient texture mapping process. It would be preferable if such a process could somehow minimize the amount of bandwidth needed for the texture mapping process. It would also be highly preferable if such a process could be implemented with fewer chips and without

redundant texture storage.

SUMMARY OF THE INVENTION

The present invention pertains to a semiconductor chip used in computer systems for performing texture mapping for graphics applications. Various texture maps are input to the semiconductor chip. These texture maps are stored in a main memory (e.g., DRAM array). The data for recent texture maps are cached in order to accelerate the reading and writing of texels. A memory controller controls the data transfers between the main memory and the cache. Also included within the same semiconductor chip is one or more interpolators. These interpolators produce an output texel by interpolating from the textures stored in memory. The interpolated texel value is output by the semiconductor chip for display. By integrating the texture memory with the integrators onto the same substrate, the transmission bandwidth as well as the redundant storage of texture maps in a multi-processor environment can be minimized.

In the currently preferred embodiment, when the DRAM row is accessed, its contents can be transferred to the cache while simultaneously accessing the cache and cycling the DRAM array again. In other words, memory access cycles can be performed in parallel to reading from a cache. Furthermore, one embodiment of the present invention incorporates fault

prediction, whereby future data is retrieved before it is currently needed. In this manner, data transfers can continue without incurring extra delays. In another embodiment, the interpolators are comprised of a number of modular slices which can be stacked to accommodate different data widths as desired.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

Figure 1 shows a block diagram of a computer graphics system upon which the present invention may be practiced.

Figure 2 shows a block diagram of a texture random access memory (TRAM).

Figure 3 shows a block diagram of the TRAM memory organization.

Figure 4 shows a block diagram of two banks of level of detail (LOD) memory.

Figure 5 shows quad DRAM and cache memories.

Figure 6 shows a circuit diagram of the TRAM resampling interconnection network.

Figure 7 shows a block diagram of a mip-map generator.

Figure 8 shows an example of a nearest-neighbor interpolation.

Figure 9 shows an example of a bi-linear interpolation.

Figure 10 shows a circuit diagram of an 8-bit bi-linear interpolator.

Figure 11 shows three examples of different configurations that can be used for the interpolator chains.

Figure 12 shows a basic 4-bit input interpolation slice.

Figure 13 shows an 8-bit input interpolator, which is formed by coupling together two 4-bit slices.

Figure 14 shows a basic 6-bit input interpolator slice.

DETAILED DESCRIPTION

An apparatus and method for integrating texture memory and interpolation logic for performing texture mapping in a computer display system is described. In the following description, for purposes of explanation, numerous specific details are set forth, such as memory sizes, cache architecture, interpolation schemes, etc., in order to provide a thorough understanding of the present invention. It will be obvious, however, to one skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

Overview of a Computer System Upon Which
the Present Invention May Be Implemented

Referring to Figure 1, a computer graphics system upon which the present invention may be practiced is shown as 100. System 100 can include any computer controlled graphics systems for generating complex or three-dimensional images, such as the IRIS™ family of computers manufactured by Silicon Graphics, Inc. of Mountain View, CA. Computer system 100 comprises a bus or other communication means 101 for communicating

information, and a processing means 102 coupled with bus 101 for processing information. System 100 further comprises a random access memory (RAM) or other dynamic storage device 104 (referred to as main memory), coupled to bus 101 for storing information and instructions to be executed by processor 102. Main memory 104 also may be used for storing temporary variables or other intermediate information during execution of instructions by processor 102. Computer system 100 also comprises a read only memory (ROM) and/or other static storage device 106 coupled to bus 101 for storing static information and instructions for processor 102. Data storage device 107 is coupled to bus 101 for storing information and instructions.

Also coupled to bus 101 is a graphics subsystem 111. Processor 102 provides the graphics subsystem 111 with graphics data such as drawing commands, coordinate vertex data, and other data related to an object's geometric position, color, and surface parameters. The object data is processed by graphics subsystem 111 in the following four pipelined stages: geometry subsystem, scan conversion subsystem, raster subsystem, and a display subsystem. The geometry subsystem converts the graphical data from processor 102 into a screen coordinate system. The scan conversion subsystem then generates pixel data based on the primitives (e.g., points, lines, polygons, and meshes) from the

geometry subsystem. The pixel data is sent to the raster subsystem, whereupon z-buffering, blending, texturing, and anti-aliasing functions are performed. The resulting pixel values are stored in frame buffer 109. The display subsystem reads the frame buffer 109 and displays the image on display monitor 121.

Furthermore, a data storage device 107 such as a magnetic disk or optical disk and its corresponding disk drive can be coupled to computer system 100. Computer system 100 can also be coupled via bus 101 to a display device 121, such as a cathode ray tube (CRT), for displaying information to a computer user. An alphanumeric input device 122, including alphanumeric and other keys, is typically coupled to bus 101 for communicating information and command selections to processor 102. Another type of user input device is cursor control 123, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 102 and for controlling cursor movement on display 121. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), which allows the device to specify positions in a plane.

Another device which may be coupled to bus 101 is hard copy device 124 which may be used for printing instructions,

data, or other information on a medium such as paper, film, or similar types of media. Additionally, computer system 100 can be coupled to a device for sound recording and/or playback 125, such as an audio digitizer coupled to a microphone for recording information. Further, the device may include a speaker which is coupled to a digital to analog (D/A) converter for playing back the digitized sounds.

Custom Texture Random Access Memory

In the present invention, the texture memory, interpolation and resampling logic associated with texture mapping is integrated onto a single chip, called the custom texture random access memory (TRAM). In other words, the TRAM circuits (e.g., texture memory such as DRAM and cache, interpolation logic, resampling logic, etc.) all reside on a same substrate of a chip. This TRAM design offers several advantages. First, by interpolating on chip, the number of data output pins as well as support chips that are need is minimized. For example, each TRAM, for an 8-bit case, only needs an 8-bit data output. The results are sent directly to the graphics rasterizer. Furthermore, by using the densest dynamic random access memory (DRAM) technology, multiple DRAMs can be created on-chip. This provides texel address independence for performing mip map

resampling, along with large texture capacity without redundant texture storage. And by caching the DRAM rows, DRAM access are run in parallel with resampling. In addition, the address and data are pipelined for increased throughput.

Figure 2 shows a block diagram of the TRAM 200. The input address and data on line 201 are input to the mip map generator 202 and the memory control and address interface 203. Once the mip maps have been generated, it is buffered by input buffer 204 before being stored by the DRAM array 205. In addition, read and cache addresses are input to the memory control and address interface 203 on lines 206 and 207. The memory control and address interface 203 controls the DRAM array 205 and the read cache 208. Sub-texel and LOD data are used in conjunction with the mip maps contained in the read cache 208 by interpolator 209. The interpolator 209 resamples the texture image to produce the output samples, which are sent to the graphics rasterizer. These blocks are described in greater detail below.

Figure 3 shows a block diagram of the TRAM memory organization. Basically, the TRAM memory is comprised of the DRAM array and the static random access memory (SRAM) cache, which are divided into two groups of parallel LOD memories 301 and 302. Each group allows for a 2x2 texel area from each of the

LOD's in order to perform tri-linear resampling. Each of these two groups has four independent "quads" 303-306 and 307-310 of memory. Dividing the LOD memory into quads allows for texture traversal for minimizing cache thrashing at the image boundaries. A double buffered load port is implemented so that the loading buffer can occur in parallel with the DRAM write cycle.

In the currently preferred embodiment, the DRAM and the I/O caches are separated into the four quads. The effect is that as a span or line is traversed during a resampling operation, the I/O caches are accessed up to four 8-bit texels from a single quad to create the sample for each LOD. These samples are then blended/muxed by the interpolate/ mux block 311 for output to the graphics rasterizer.

Note that the four quads of each group have independent addressing as shown by the four buses 312-315. Hence, each of the quads can walk in the scan direction, prefetching data that will be needed along the scan path. The graphics rasterizer monitors the page address in a pipeline and pre-determines whether a page fault will occur. It will then start the access to the necessary data to avoid the fault (i.e., fault prediction). By the time the address reaches the end of the pipe, where it is used

to resample from the TRAM cache, the data is already available from the DRAM array.

Figure 4 shows a block diagram of two "banks" 401-402 of LOD memory. These two banks represent 1 Meg x 8 quads of memory. Bank 401 is comprised of four 64x1 byte write masks 403-406; four 64x8 out/in 512x1 in/out I/O caches 407-410; a 1Kx256 even rows DRAM 411; and a 1Kx256 odd rows DRAM 412. Similarly, the second bank 402 is comprised of blocks identical to the first bank 401.

These two banks allow for a "ping-pong" mode, wherein load and resampling performance are decoupled (i.e., full performance on both ports). This feature is provided for 4-bit and 8-bit textures. In the currently preferred embodiment, flexibility is provided in the configurations for each LOD. This allows for either ping-ponged 4-bit or 8-bit banks, spatially larger 4-bit or 8-bit banks, or the combination of the two banks to increase the texel depth (e.g., 12 bits). For non-mip mapped applications, up to a 1Kx1Kx8-bit texture map can be stored.

Figure 5 shows the different ways in which the DRAM and cache memories can be organized for a given quad. The 16x16 array 501 represents one of the 128x8 bytes in each 1K DRAM. Array 501 can be a page in a quad belonging to either of the two

banks. Row 502 represents an even 1K DRAM, while row 503 represents an odd 1K DRAM.

Configuration 504 shows the two banks 505 and 506 used in the ping-pong mode as described above. In one embodiment, banks 505 and 506 have a 16x16x8-bit size. In another embodiment, these two banks can be configured to a 16x32x4-bit size. Note that in the 4 and 8-bit modes, parallelized loading and resampling can be performed to the same TRAM for ping-pong mode. This allows independent operation and utilizes resampling resources (e.g., column decode, I/O cache, etc.).

Alternatively, the two banks 505 and 506 can be placed side-by-side in series as shown in configuration 507. This configuration 507 provides for a spatially larger contiguous memory. In yet another embodiment, the two banks are placed in a paralleled configuration 508. Configuration 508 allows for a deeper memory for storing texel depths of up to 16 bits.

When sampling is performed at the middle of an image, data is needed from both banks. However, the retrieved data is found in complementary quads, so that all of the data is available simultaneously. With non-mip mapped and non-enabled borders, the data at these boundaries will come from the other group. When either mip mapped or when borders are enabled, each bank will hold this border data in its border store. When

mip mapping is enabled or borders are enabled, operation will be consistent (i.e., data at borders will come from the same group).

The memory control and address interface 203 of Figure 2 controls the DRAM arrays and also controls the sharing of the DRAM array for non-ping pong modes, as the load and resampling ports require the same DRAM resource. This interface is used in swapping of the two banks in ping-pong modes. Interface 203 interfaces the addresses from the graphics rasterizer to the row addresses for each quad's DRAMs. This includes double buffered page address registers for each DRAM for parallel DRAM access and page address register loading. In addition, the memory control and address interface stores and increments the load port DRAM page address for single image or mip map generation.

An interconnection network connects the quads to the interpolators and the data input logic. In the interpolation process, the four quads of a given group are multiplexed. These results are then made available to each interpolator (i.e., they are multiplexed at its inputs). The two levels of muxing allows the entire array to be used as one large image or as two independent LODs. Border data originates from the same group, but is stored in an opposite quad from the image edge data (e.g., horizontal edges are stored in AB<->CD; vertical edges are stored in AC<->BD).

In either case, each LOD is split between the two groups. LODs which are smaller than 16x16 can wrap because: 1) the other horizontal half is in the other group so that quad "a" and quad "b" data from the same 16x16 area is not needed, hence, they will always be contiguous and 2) the vertical quad "ab" and quad "cd" addresses are independent so that any two lines can be read from the same 16x16 cache.

Figure 6 is a circuit diagram of the TRAM resampling interconnection network. Two groups 601 and 602 of memory, each having four quads "a's-d's" 603-606 and 607-610 are coupled to sense amplifiers 611. The sense amplifiers 611 are used to amplify the weak signals from the two group's DRAM data lines. A giant latch 612 is used to pipeline the data after the first level of muxing. Four 2:1 multiplexers 621-624 are used to allow a portion of or each LOD to be stored in each group. Thus, the coefficients to the interpolators 625 and 626 provides for the muxing of two inputs. The outputs from interpolators 625 and 626 are in turn interpolated by interpolator 627. A similar structure exists for the other LOD. A final interpolator 629 is used to interpolate the results from both LODs as output from the interpolators 627 and 628.

For loading the TRAM, two LODs are simultaneously sent, when using the mip map generator, to either group 601 or 602

and to any texel location within that group. Since the 4:1 multiplexers 630-637 are comprised of bi-directional transmission gates, they can either be used for resampling or writing texels. When storing borders, as with resampling, the data is stored to the appropriate memories. When borders or mip mapping is enabled, the center border is stored twice, once in each group. A number of drivers 613-620 are used in loading data to the DRAMs. Two drivers are associated with each input line so that the two bytes of data can separately be enabled.

The mip map generator 202 as shown in Figure 2 can create an entire mip map for complete images. The load input of the TRAM allow the texture to be loaded in a rectilinear fashion. For tiled inputs, the mip map generator can filter the input tile down to a minimum of one texel. The mip map generator can be bypassed if it is necessary to explicitly load the mip map levels as would be the case if a user-defined filter is to be applied to create the mip maps (e.g., CPU generated).

In the currently preferred embodiment, the mip map generator performs a 2x2 box filter (i.e., averaging filter) to produce each texel for the next rougher mip map level. A box filter eliminates most of the high frequencies that would cause aliasing, but it is not a very sharp filter. The mip map generator handles 4, 8, and 12-bit images as well as four 4-bit and two 8-

bit images for a single TRAM configuration. With multiple TRAMs, the mip maps are generated and stored in parallel.

Figure 7 shows a block diagram of a mip map generator.

The input stage 701 performs the horizontal filtering on the fine LOD level. The second stage 702 takes the results from the input stage 701 and generates an output and temporary results for LODs further down the mip map. The adders 703 and 704 can be used as four 4-bit adders in parallel, two 8-bit adders, or one 12-bit adder, depending on the selected texel mode. These two adders 703 and 704 basically averages, with truncation, two values by adding then dividing by two. In one instance, two 4-bit numbers are input to produce a 4-bit result. In other instances, 8 and 12-bit numbers are input to the adders to produce 8 and 12-bit results. The width of these adders need only be as wide as the data. Since 4-bit (1 to 4 components), 8-bit (1 or 2 components), or 12-bit data are supported, four 4-bit adder slices are implemented. These four 4-bit adder slices can be combined to create 8 or 12-bit adders. A number of line stores 705-710 are used for temporarily storing interim processing results. Line stores 711 and 712 are used to temporarily store the outputs from the mip map generator. The address generator 713 and the 512x16 line RAM 714 are used to read/write addresses and data to/from the mip map generator.

All components are input, processed and stored to DRAM in parallel. Hence, the only difference is the carry circuitry at the adders. The RAM allows for larger line length, so the two and four TRAM configurations allow for larger image lines which goes along with increases in texture RAM capacity.

In the currently preferred embodiment, mip maps and borders are stored in TRAM in a contiguous array of TRAM blocks. The size of these blocks are either 8Kbits or 16Kbits, depending on whether banks 1 and 2, in each LOD's memory, are used together or not. The basic unit of memory allocation is 8Kbits. This allows for the following combined page/cache sizes: 1) 32x32x12 bits, where banks 1 and 2 always used together -- 16Kbits; 2) 32x32x8 bits or 32x64x4 bits, with two banks for ping-ponging -- 8Kbits each; and 3) 64x32x8 bits or 64x64x4 bits, with one bank for large single images --16Kbits.

There is a BASE offset register that is used to specify the starting address of the image data. Hence, all mip map and border offsets are relative to the base. These offsets do not change if the image were put in a different part of TRAM at some later time. To provide for more efficient utilization of TRAM when borders are being used or when many small textures are bound, borders and small mip maps are packed whenever possible into the same page. Each mip map and each border for

each mip map has a table entry for the relative offset from the BASE address to the start of that data. The offsets include sub-block address bits so that mip maps and borders do not have to be aligned to page boundaries. For example, the 16x16, 8x8, 4x4, 2x2, and 1x1 map levels can be stored together in one memory page. In addition, the borders can be similarly packed.

The mip map interpolators 209 of Figure 2 resamples the texture image to produce the output samples. The goal is to approximate as closely as possible what the texture image pixel value would have been, at the location requested, if a continuous texture image were given. Low pass filtering is implemented to minimize aliasing. The interpolation circuit generates the approximation while the mip map generator handle the low pass filtering.

The hardware performs a "nearest-neighbor/point sample" (NN) or bi-linear interpolation on each mip map level. For NN resampling (sometimes called "point" or "nearest" mode), if the (S,T) coordinate falls within a texel's boundary after truncation, that level's sample is assigned the intensity of the corresponding texel. Figure 8 shows an example of nearest neighbor resampling. Six texels 801-806 having intensities of 40, 50, 60, 30, 45, and 55, respectively, are shown. Given a texel coordinate of (S,T)=(1.75,1.25), its NN output corresponds to texel 802, and

an intensity of 50 will be used. This nearest neighbor sampling scheme produces a blocky appearance when the image is magnified. It also produces more edge aliasing when an edge in the image is rotated with respect to the screen pixels (i.e., jaggies). Consequently, bi-linear filtering is used to "smooth out" these effects.

In bi-linear resampling, a two dimensional interpolation is performed. The intensity of the texels having the four nearest centers are used. Figure 9 illustrates an example of bi-linear resampling. Four texels 901-904 are shown respectively having four intensities I0-I3. The intensity for point 905 is given by the following equation: $I = Ct((Cs(I1-I0)+I0)-(Cs(I3-I2)+I2))+(Cs(I3-I2)+I2)$, where Cs and Ct are the offsets between the (S,T) coordinates of texel 901. Hence, the output intensity resulting from the bi-linear resampling is a function of the four intensities I0-I3.

Figure 10 shows a bi-linear resampling circuit diagram. The ABS blocks 1001 and 1002 output the absolute values of the S, T, or LOD value to compensate for the fact that whether the sample location is on the right or left of a texel, the texel goes to the same input of the interpolator. These absolute values make the coefficient track the texel data. The output from ABS block 1001 gives the value of Ct, and the output from ABS block 1002

gives the value of C_s . Subtractor 1003 performs a subtraction function on the intensities I_0 and I_1 . Multiplier 1004 multiplies this result by C_s . Next, adder 1005 adds I_0 , thereby resulting in an output of $C_s(I_1 - I_0) + I_0$. In similar fashion, the output from adder 1006 is $C_s(I_3 - I_2) + I_2$. Subtractor 1007 subtracts the output of adder 1006 from the output of adder 1005. This gets multiplied by C_t by multiplier 1008. Lastly, the output from multiplier 1008 is added to the output from adder 1006, which results in the bi-linear sampled intensity.

Once these resamplings are completed, the hardware either selects between one of these schemes or linearly interpolates between them to create the output sample. These results then get interpolated in the mip map generator block. Both nearest neighbor and bi-linear filtering still apply when using a mip map. These filters are performed on one or two map levels, depending on the filter selected. Each TRAM interpolates 1 (4 TRAM configuration), 2 (2 TRAM configuration), or 4 (single TRAM configuration) components. For other types of formats, the texel color values are truncated to fit into the available texel depth.

It has been determined that to maintain error at less than +/- 1.0 LSB at the LOD interpolator output, the two fraction bits kept and rounding is performed using the third, at the output of the multiply is necessary. Maintaining two fraction bits builds

error to +/- 3/4 in the currently preferred embodiment, each linear interpolator is "sliced" to allow for 4, 8 or 12-bit data. The first stage of interpolators only need to handle 4-bit input data. The second and third stages are comprised of a slice that handles 6-bit input data.

Figure 11 shows three examples of different configurations that can be used for the interpolator chains. In the first configuration, four of the 4-bit unit 1101 can be stacked to perform 16-bit tri-linear interpolations. In the second configuration, a pair of the 8-bit unit 1102 can be stacked to perform the 16-bit tri-linear interpolations. In the third configuration, a single 12-bit unit 1103 is used to perform the interpolations. A number of multiplexers are programmed to couple the various units in order to achieve the desired configuration. In this manner, the same hardware can be utilized with minimal extra gate count to attain the flexibility in data width. In some applications, greater precision and resolution are desirable. Whereas, in other applications, speed and cost are of greater importance.

In the currently preferred embodiment, the units 1101-1103 are comprised of an S-stage, a T-stage, and an LOD-stage. The S-stages are comprised of 4-bit slices 1104. The T-stages and LOD stages are comprised of 6-bit slices 1105. These slices

are designed to make the interpolators modular. Figure 12 shows a basic 4-bit input interpolation slice. The data to the muxes comes from the subtractor (not shown). The most significant (MS) bits are either sign extension (if the slice is being used as the MS slice) or are the low order bits from the next more significant slice (when not the most significant slice). Since each slice handles both a sign bit and its extension, as well as fraction bits, these capabilities are disabled or bypassed as appropriate for the position of the slice in the larger word. For example, for 8-bit texels, the MS slice would handle the sign bit (and no fraction bits) and the least significant (LS) slice would handle the fraction bits (but no sign). The slice configuration muxes separately enable or disable the slices MS portion (sign extension) or LS portion (fraction). Figure 13 shows an 8-bit input interpolator, which is formed by coupling together two 4-bit slices. Figure 14 shows a basic 6-bit input interpolator slice.

The adders and subtractors are split into carry look-ahead adders (CLAs) that have looked ahead across the slices to allow for the desired combinations. The multipliers use a modified Booth (e.g., bit pair recoding) followed by a Wallace tree, whereby the input from the subtractor will be combined in the final stage of the Wallace tree. The multiplier is constructed such

that the entire interpolator slice may be combined to form the higher precision versions. The multiply is $I * (Sa[0:5] + Sb[0:5] + Sc[0:5] + Sd[0:5])$. This allows the individual nibbles to be interpreted separately or shifted appropriately and added, to create either 4-bit, 8-bit or 12-bit results (e.g., for the 6-bit slices). "I" is the 6-bit interpolation value (multiplier in this case). This basic multiplier is thus 7x6. Because the multiplier ("I") is the same across each group of four slices, the same recode logic can be used across groups of four slices.

In the currently preferred embodiment, instead of combining the results after they have individually been computed, an approach is taken that uses multiplier slices that can be connected, in parallel, for the wider precision. This prevents requiring full precision (6+4=10 bits) in each slice and eliminates the final stage of adders that would be required.

Thus, an integrated texture memory and interpolation logic device used in performing texture mapping in a computer display system is disclosed.

CLAMS**What is claimed is:**

1. In a computer system, a semiconductor chip for performing texture mapping, said semiconductor chip comprising:
 - an input for inputting textures to said semiconductor chip;
 - a main memory coupled to said input means for storing said textures;
 - a cache memory coupled to said main memory for storing a recently used texture;
 - a memory controller coupled to said main memory and said cache memory for controlling data transfers between said main memory and said cache memory;
 - an interpolator coupled to said memory for producing an output texel by interpolating from said recently used texture stored in said cache memory;
 - an output coupled to said interpolator for outputting said output texel, wherein said input, said main memory, said cache memory, said memory controller, and said interpolator reside on a same substrate.

2. The semiconductor chip of Claim 1, wherein said memory controller performs memory access cycles to said main memory in parallel with read cycles to said cache memory.

3. The semiconductor chip of Claim 2, wherein said main memory and said cache memory are divided into at least two independently addressable banks to decouple loading and resampling operations.

4. The semiconductor chip of Claim 3 further comprising a means for predicting faults, said means for predicting faults residing on said substrate.

5. The semiconductor chip of Claim 1, wherein said interpolator is comprised of a plurality of modular slices which can be stacked to generate interpolations for different data widths.

6. The semiconductor chip of Claim 5, wherein said interpolator performs a nearest neighbor selection of said texture to produce said output texel.

7. The semiconductor chip of Claim 5, wherein said interpolator performs a bi-linear interpolation of said texture to produce said output texel.

8. The semiconductor chip of Claim 1 further comprising a filter coupled to said input for filtering said textures prior to storage in said main memory.

9. The semiconductor chip of Claim 8, wherein said filter is comprised of a 2x2 box filter.

10. In a computer system, a method of performing texture mapping, said method comprising the steps of:

inputting textures to a semiconductor chip;
storing said textures in a main memory of said semiconductor chip;

storing a recently used texture in a cache memory of said semiconductor chip;

controlling data transfers between said main memory and said cache memory;

producing an output texel by implementing an interpolator and interpolating from said recently used texture stored in said

cache memory, wherein said interpolator resides on said semiconductor chip;
outputting said output texel from said semiconductor chip.

11. The method of Claim 10 further comprising the step of performing memory access cycles to said main memory in parallel with read cycles to said cache memory.

12. The method of Claim 11 further comprising the steps of dividing said main memory and said cache memory into at least two independently addressable banks to decouple loading and resampling operations.

13. The method of Claim 12 further comprising the step of predicting faults by determining data that will be required and fetching said data before said data is required.

14. The method of Claim 10 further comprising the step of stacking a plurality of modular slices to generate interpolations for different data widths.

15. The method of Claim 14 further comprising the step of said interpolator performing a nearest neighbor selection of said texture to produce said output texel.


16. The method of Claim 14 further comprising the step of said interpolator performing a bi-linear interpolation of said texture to produce said output texel.

17. The method of Claim 10 further comprising the step of filtering said textures prior to storage in said main memory.



18. The method of Claim 17, wherein said filtering step is performed by utilizing a 2x2 box filter.

19. A single semiconductor chip comprising a cache memory for storing texture information and an interpolator for generating a texel by interpolation from said texture information stored in said cache memory, wherein said cache memory and said interpolator both reside on said single semiconductor chip.

1/11

FIG 1

2/11

FIG — 2**FIG — 3**

3/11




FIG 4


FIG 5



16



5/11



6/11

卷之七





FIG — **X**



7/11

FIG - 9






FIG _10




8/11

FIG. III


9/11

FIGURE 12




10/11

FIG. 13



11/11

FIG 14



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US95/02853

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06F 12/00

US CL : 395/ 164

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/119, 125, 126, 128, 129, 130, 162, 164, 166

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS, DIALOG

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US, A, 5,222,205 (LARSON ET AL) 22 JUNE 1993, col. 3, lines 41-68, col. 7 lines 23-45.	8, 9, 17, 18
X	Williams, Tom: "80860 May Force Rethinking of Graphics System Architectures", Computer Design, v28, n9, p43(3), 01 July 1989.	1, 4, 10, 13, 19
Y		----- 2, 3, 5-7, 11, 12, 14-17

 Further documents are listed in the continuation of Box C. See patent family annex.

- Special categories of cited documents:
 - *A* document defining the general state of the art which is not considered to be part of particular relevance
 - *E* earlier document published on or after the international filing date
 - *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
 - *O* document referring to an oral disclosure, use, exhibition or other means
 - *P* document published prior to the international filing date but later than the priority date claimed
- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "Z" document member of the same patent family

Date of the actual completion of the international search

19 JUNE 1995

Date of mailing of the international search report

04 AUG 1995

Name and mailing address of the ISA/US
 Commissioner of Patents and Trademarks
 Box PCT
 Washington, D.C. 20231
 Facsimile No. (703) 305-3230

Authorized officer
 MARK POWELL *Joni Hill*
 Telephone No. (703) 305-9703

**.This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.